

IN THE CLAIMS:

Please amend the claims as follows:

1-28. (Canceled).

29. (Currently Amended) A computer-implemented method, implemented on a computer, for automatically providing a marketing strategy to address at least one specified merchant objective, the objective corresponding to a specified time period and a specified budget, the strategy being implemented across at least one marketing channel, the strategy including at least one initiative, the computer-implemented method comprising:

inputting to a computer, by a merchant, said at least one specified merchant objective, said at least one specified merchant objective including said specified time period and said specified budget as constraints;

generating, by said computer, a plurality of possible marketing strategies, each of said plurality of possible marketing strategies comprising a set of initiatives, which are deployed together in a given sequence for said specified time period,

wherein said initiatives include any of bundling of products, cross-sells, up-sells, coupons, discounts, promotions, advertisements, surveys, and customer feedback;

determining, by said computer, an optimal marketing strategy from said plurality of possible marketing strategies, each of said plurality of possible marketing strategies corresponding to a policy comprising a sequence of various actions taken at different states encountered during said specified time period,

wherein each of said different states corresponds to a set of variables including any of customer profile, purchase frequency, and monetary value of purchase, associated with a customer at a time in said specified time period,

wherein said various actions include at least one randomized action;

wherein said various actions are constrained by a choice of a marketing channel,

wherein said various actions are deployed across said different states encountered during said specified time period to provide a plurality of policies,

wherein each of said plurality of policies is evaluated in a context of a reinforcement learning algorithm, in which values of each of said plurality of policies corresponds to a vector of total expected rewards, said total expected rewards comprising a sum of rewards corresponding to a monetary value for each of said various actions deployed across each of said different states during said specified time period,

wherein said determining, by said computer, of said optimal marketing strategy comprises determining an optimal policy for each state during said specified time period based on past data, wherein said determining of said optimal policy comprises:

evaluating, in order, said each state for said specified time period across each of said plurality of policies, corresponding to said sum of rewards; and

identifying said optimal policy, associated with said optimal marketing strategy, by a maximal value representing a maximal total expected reward for said optimal policy; and

outputting, by said computer, to said merchant, the optimal marketing strategy.

30. (Currently Amended) The computer-implemented method according to claim 29, wherein generating a plurality of possible marketing strategies comprises:

selecting at least one initiative that enables an addressing of the at least one specified merchant objective;

determining sequences in which selected initiatives are deployed, if more than one initiative is selected; and

combining the selected initiatives in the determined sequences to obtain the plurality of possible marketing strategies.

31. (Currently Amended) The computer-implemented method according to claim 30, further comprising varying parameters of initiatives to generate new initiatives, corresponding to new states during said specified time period.

32. (Currently Amended) The computer-implemented method according to claim 30, further comprising varying deployment time of said initiatives.

33. (Currently Amended) The computer-implemented method according to claim 29, wherein the determining an optimal marketing strategy further comprises, after said determining of said optimal policy for each state based on past data:

identifying a state of a customer, a customer visiting a merchant, or a customer being selected from a database of customers; and

identifying an optimal marketing strategy using the state of the customer, the identified optimal policy, and constraints corresponding to marketing channels.

34-35. (Canceled).

36. (Currently Amended) The computer-implemented method according to claim 29, wherein evaluating, in order, said each state for said specified time period across each of said plurality of policies comprises:

computing transition probabilities from a given state to another state;

computing a value of expected immediate reward corresponding to said given state;

computing a discounted expected value of a resulting state; and

computing a sum of expected immediate rewards and a total discounted expected value corresponding to a sum of states.

37. (Canceled).

38. (Currently Amended) The computer-implemented method according to claim 33, wherein the identifying an optimal marketing strategy comprises:

identifying the optimal policy for an identified customer state;

assigning customer's preferences for marketing channels, cost, and effectiveness of different marketing channels, and the specified budget as constraints;

determining an optimal feasible policy based on the identified optimal policy and constraints corresponding to marketing channels; and

determining an optimal feasible marketing strategy from the optimal feasible policy.

39. (Currently Amended) The computer-implemented method according to claim 38, wherein the determining an optimal feasible policy based on constraints corresponding to marketing channels comprises mapping the optimal policy uniquely to a closest feasible optimal policy based on the constraints, if the constraints are not satisfied by the optimal policy.

40-42. (Canceled).

43. (Currently Amended) A computer system that implements a computer-implemented method for automatically providing a marketing strategy to address at least one specified merchant objective, the objective corresponding to a specified time period and a specified budget, the strategy being implemented across at least one marketing channel, the strategy including at least one initiative, the computer system comprising:

a memory for storing said at least one specified merchant objective, which is inputted by a merchant via an input device and a communication bus, said at least one specified merchant objective including said specified time period and said specified budget as constraints; and
a microprocessor configured to:

generate a plurality of possible marketing strategies, each of said plurality of possible marketing strategies comprising a set of initiatives, which are deployed together in a given sequence for said specified time period,

wherein said initiatives include any of bundling of products, cross-sells, up-sells, coupons, discounts, promotions, advertisements, surveys, and customer feedback;

determine an optimal marketing strategy from said plurality of possible marketing strategies, each of said plurality of possible marketing strategies corresponding to a policy

comprising a sequence of various actions taken at different states encountered during said specified time period,

wherein each of said different states corresponds to a set of variables including any of customer profile, purchase frequency, and monetary value of purchase, associated with a customer at a time in said specified time period,

wherein said various actions include at least one randomized action;

wherein said various actions are constrained by a choice of a marketing channel and said specified budget,

wherein said various actions are deployed across said different states encountered during said specified time period to provide a plurality of policies,

wherein each of said plurality of policies is evaluated in a context of a reinforcement learning algorithm, in which values of each of said plurality of policies corresponds to a vector of total expected rewards, said total expected rewards comprising a sum of rewards corresponding to a monetary value for each of said various actions deployed across each of said different states during said specified time period,

wherein said determining of said optimal marketing strategy comprises determining an optimal policy for each state during said specified time period based on past data, wherein said determining of said optimal policy comprises:

evaluating, in order, said each state for said specified time period across each of said plurality of policies, corresponding to said sum of rewards; and

identifying said optimal policy, associated with said optimal marketing strategy, by a maximal value representing a maximal total expected reward for said optimal policy; and

~~outputting~~ output, to said merchant, the optimal marketing strategy.

44. (Currently Amended) The computer system according to claim 43, wherein said generating a plurality of possible marketing strategies comprises:

selecting at least one initiative that enables an addressing of the at least one specified merchant objective;

determining sequences in which selected initiatives are deployed, when more than one initiative is selected; and

for combining the selected initiatives in the determined sequences to obtain the plurality of possible marketing strategies.

45. (Currently Amended) The computer system according to claim 43, wherein the identifying an optimal marketing strategy comprises:

identifying a state of a customer, a customer visiting a merchant, or a customer being selected from a database of customers;

identifying an optimal policy for an identified customer state;

assigning customer's preferences for marketing channels, cost, and effectiveness of different marketing channels, and the specified budget as constraints;

determining an optimal feasible policy based on constraints corresponding to marketing channels; and

determining an optimal feasible marketing strategy from the optimal feasible policy.

46. (Canceled).

47. (Currently Amended) The computer system according to claim 45, wherein determining an optimal feasible policy based on constraints corresponding to marketing channels comprises mapping the optimal policy uniquely to a closest feasible optimal policy based on the constraints, if the constraints are not satisfied by the optimal policy.

48. (Canceled).

49. (Currently Amended) A computer program storage device readable by computer, tangibly embodying a computer program of instructions executable by the computer to perform a computer-implemented method for automatically providing a marketing strategy to address at least one specified merchant objective, the objective corresponding to a specified time period and

a specified budget, the strategy being implemented across at least one marketing channel, the strategy including at least one initiative, the computer-implemented method comprising:

inputting to a computer, by a merchant, said at least one specified merchant objective, said at least one specified merchant objective including said specified time period and said specified budget as constraints;

generating, by said computer, a plurality of possible marketing strategies, each of said plurality of possible marketing strategies comprising a set of initiatives, which are deployed together in a given sequence for said specified time period,

wherein said initiatives include any of bundling of products, cross-sells, up-sells, coupons, discounts, promotions, advertisements, surveys, and customer feedback;

determining, by said computer, an optimal marketing strategy from said plurality of possible marketing strategies, each of said plurality of possible marketing strategies corresponding to a policy comprising a sequence of various actions taken at different states encountered during said specified time period,

wherein each of said different states corresponds to a set of variables including any of customer profile, purchase frequency, and monetary value of purchase, associated with a customer at a time in said specified time period,

wherein said various actions include at least one randomized action;

wherein said various actions are constrained by a choice of a marketing channel and said specified budget,

wherein said various actions are deployed across said different states encountered during said specified time period to provide a plurality of policies,

wherein each of said plurality of policies is evaluated in a context of a reinforcement learning algorithm, in which values of each of said plurality of policies corresponds to a vector of total expected rewards, said total expected rewards comprising a sum of rewards corresponding to a monetary value for each of said various actions deployed across each of said different states during said specified time period,

wherein said determining of said optimal marketing strategy comprises determining an optimal policy for each state during said specified time period based on past data, wherein said determining of said optimal policy comprises:

evaluating, in order, said each state for said specified time period across each of said plurality of policies, corresponding to said sum of rewards; and

identifying said optimal policy, associated with said optimal marketing strategy, by a maximal value representing a maximal total expected reward for said optimal policy; and

outputting, by said computer, to said merchant, the optimal marketing strategy.

50. (Currently Amended) The computer program storage device according to claim 49, wherein the generating a plurality of possible marketing strategies comprises:

selecting at least one initiative that enables an addressing of the at least one specified merchant objective;

determining sequences in which selected initiatives are deployed, when more than one initiative is selected; and

combining the selected initiatives in the determined sequences to obtain the plurality of possible marketing strategies.

51. (Currently Amended) The computer program storage device according to claim 49, wherein the determining an optimal marketing strategy comprises:

identifying a state of a customer, a customer visiting a merchant, or a customer being selected from a database of customers;

identifying an optimal policy for an identified customer state;

assigning customer's preferences for marketing channels, cost, and effectiveness of different marketing channels, and the specified budget as constraints;

determining an optimal feasible policy based on constraints corresponding to marketing channels; and

determining an optimal feasible marketing strategy from the optimal feasible policy.

52-55. (Canceled).

56. (Currently Amended) A computer-implemented method, implemented on a computer, for automatically providing a marketing strategy to address at least one specified merchant objective, the objective corresponding to a specified time period and a specified budget, the strategy being implemented across at least one marketing channel, the strategy including at least one initiative, the computer-implemented method comprising

inputting to a computer, by a merchant, said at least one specified merchant objective, said at least one specified merchant objective including said specified time period and said specified budget as constraints;

generating, by said computer, a plurality of possible marketing strategies, each of said plurality of possible marketing strategies comprising a set of initiatives, which are deployed together in a given sequence for said specified time period,

wherein said initiatives include any of bundling of products, cross-sells, up-sells, coupons, discounts, promotions, advertisements, surveys, and customer feedback;

determining, by said computer, an optimal marketing strategy from said plurality of possible marketing strategies, each of said plurality of possible marketing strategies corresponding to a policy comprising a sequence of various actions taken at different states encountered during said specified time period,

wherein each of said different states corresponds to a set of variables including any of customer profile, purchase frequency, and monetary value of purchase, associated with a customer at a time in said specified time period,

wherein said various actions include at least one randomized action;

wherein said various actions are constrained by a choice of a marketing channel and said specified budget,

wherein said various actions are deployed across said different states encountered during said specified time period to provide a plurality of policies,

wherein each of said plurality of policies is evaluated in a context of a reinforcement learning algorithm, in which values of each of said plurality of policies corresponds to a vector of total expected rewards, said total expected rewards comprising a sum of rewards corresponding to a monetary value for each of said various actions deployed across each of said different states during said specified time period,

wherein said determining, by said computer, of said optimal marketing strategy comprises determining an optimal policy for each state during said specified time period based on past data, wherein said determining of said optimal policy comprises:

evaluating, in order, said each state for said specified time period across each of said plurality of policies, corresponding to said sum of rewards; and

identifying said optimal policy, associated with said optimal marketing strategy, by a maximal value representing a maximal total expected reward for said optimal policy;

outputting, by said computer, to said merchant, the optimal marketing strategy;

recording customer response to the outputted optimal marketing strategy; and

updating information corresponding to a state of a customer based on the recorded customer response.